

6 Superposition-Based Methods

Superposition is a refutationally complete calculus for first-order clauses with equality. It is an extension of both ordered resolution and Knuth/Bendix completion.

Literature (on the course web page):

Armando, Ranise, Rusinowitch: CSL 2001 paper

Lynch, Morawska: LICS 2002 paper

Bachmair, L. and Ganzinger, H.: Rewrite-based equational theorem proving with selection and simplification, *Journal of Logic and Computation* 4(3), 1994, 217–247. (available from my personal web page)

Basic Notions

1

Reduction ordering: well-founded (partial) ordering \succ on T_Σ (ground terms!) compatible with function application. That is, if $s \succ t$ then also $f(\dots, s, \dots) \succ f(\dots, t, \dots)$.

General assumption in the sequel: \succ is an arbitrary but fixed total reduction ordering. In each application we will choose a specific suitable \succ .

Extension of \succ to ground literals: (We have only equational atoms and no predicate symbols.) If $s \approx t$ is a positive ground equation, let $c(s \approx t) = (\max_\succ(s, t), 0, \min_\succ(s, t))$. For negative literals $s \not\approx t$, let $c(s \not\approx t) = (\max_\succ(s, t), 1, \min_\succ(s, t))$. Now, if L and L' are two ground literals, define $L \succ L'$ iff $c(L) > c(L')$, comparing these triples lexicographically, using \succ for terms and $1 > 0$ on “signs”.

Basic Notions

2

Extension of \succ to ground clauses: Define $C \succ D$ iff $C \succ_{ms} D$, with \succ_{ms} the multiset extension of \succ on ground literals. This gives us a total and well-founded ordering on ground clauses.

Extension of \succ to non-ground expressions: Define $E \succ E'$ iff $E\sigma \succ E'\sigma$ for all grounding substitutions σ .

On **non-ground** expressions, \succ is only **partial** in general, and it makes sense to use $\not\succeq$ and $\not\prec$ for the complement of \succ and \preceq , respectively.

Superposition Calculus \mathcal{S}

3

Let \succ be an atom ordering and S a selection function. A literal L is called **[strictly] maximal** wrt. a clause C if there exists a ground substitution σ such that for all L' in C : $L\sigma \succeq L'\sigma$ [$L\sigma \succ L'\sigma$].

$$\frac{C \vee u \approx v \quad D \vee s[u'] \approx t}{(C \vee D \vee s[v] \approx t)\sigma} \quad \text{[positive superposition]}$$

if $\sigma = \text{mgu}(u, u')$ and

- (i) u' is not a variable;
- (ii) $u\sigma \not\prec v\sigma$;
- (iii) $(u \approx v)\sigma$ is strictly maximal wrt. $C\sigma$;
- (iv) $s\sigma \not\prec t\sigma$;
- (v) $(s \approx t)\sigma$ is strictly maximal wrt. $D\sigma$;
- (vi) nothing is selected in C and D by S .

Negative Superposition

4

$$\frac{C \vee u \approx v \quad D \vee s[u'] \not\approx t}{(C \vee D \vee s[v] \not\approx t)\sigma} \quad [\text{negative superposition}]$$

if $\sigma = \text{mgu}(u, u')$ and

- (i) u' is not a variable;
- (ii) $u\sigma \not\approx v\sigma$;
- (iii) $(u \approx v)\sigma$ is strictly maximal wrt. $C\sigma$;
- (iv) $s\sigma \not\approx t\sigma$;
- (v) either $s \not\approx t$ is selected, or else nothing is selected in $D \vee s \not\approx t$ and $(s \not\approx t)\sigma$ is maximal wrt. $D\sigma$;
- (vi) nothing is selected in C .

Reflexivity and Factoring

5

$$\frac{D \vee s \not\approx t}{D\sigma} \quad [\text{reflexivity resolution}]$$

if $\sigma = \text{mgu}(s, t)$ and either $s \not\approx t$ is selected,
or else nothing is selected in $D \vee s \not\approx t$ and $(s \not\approx t)\sigma$ is maximal wrt. $D\sigma$.

$$\frac{C \vee u \approx v \vee u' \approx t}{(C \vee D \vee v \not\approx t \vee u \approx t)\sigma} \quad [\text{Equality factoring}]$$

if $\sigma = \text{mgu}(u, u')$ and

- (i) $u\sigma \not\approx v\sigma$;
- (ii) $v\sigma \not\approx t\sigma$;
- (iii) $(u \approx v)\sigma$ is maximal wrt. $C\sigma$;
- (iv) nothing is selected in C .

A Formal Notion of Redundancy

6

Ground case: Let N be a set of ground clauses and C a ground clause (not necessarily in N).

C is called **redundant** in N $:\Leftrightarrow$ there exists $C_1, \dots, C_n \in N$, $n \geq 0$:
 $C_i \prec C$ and $C_1, \dots, C_n \models C$

Redundancy for general clauses:

C is called **redundant** in N $:\Leftrightarrow C\sigma$ redundant in $G_\Sigma(N)$,
for all ground instances $C\sigma$ of C

Intuition: Redundant clauses are no minimal counterexamples for any interpretation.

NB: The same ordering \succ is used both for ordering restrictions and for redundancy. “ \models ” is validity in FOL with equality.

Examples of Redundancy

7

PROPOSITION 6.1

- If C is a tautology (i.e., $\models C$) then C is redundant in any set N .
- If $C\sigma \subset D$ then D is redundant in $N \cup \{C\}$
- $C[s\sigma]$ is redundant in any $N \cup \{C[t\sigma], s \approx t\}$, if $s\sigma \succ t\sigma$ and $C[s\sigma] \succ (s \approx t)\sigma$
- C is redundant in $N \cup \{C\sigma\}$, whenever the set of literals in $C\sigma$ is a proper subset of the set of literals in C . In that case we call $C\sigma$ a **condensment** of C .

N is called **saturated up to redundancy** (wrt. \mathcal{S}) iff

$$\mathcal{S}(N \setminus \text{Red}(N)) \subseteq N \cup \text{Red}(N),$$

where $\text{Red}(N)$ denotes the set of clauses that are redundant in N .

THEOREM 6.2 (BACHMAIR, GANZINGER 1994) Let N be saturated up to redundancy wrt. \mathcal{S} . Then $N \models \perp$ if, and only if, $\perp \in N$.

Superposition Derivations

$\triangleright_{\mathcal{S}}$ is this relation on sets (modulo variable renaming) of equational clauses, modelling both deduction and simplification:

Tautology elimination: $N \cup \{C\} \triangleright_{\mathcal{S}} N$,
if C is a tautology

Subsumption: $N \cup \{C, D\} \triangleright_{\mathcal{S}} N \cup \{C\}$,
if C strictly subsumes D

Reduction: $N \cup \{C[s\sigma], s \approx t\} \triangleright_{\mathcal{S}} N \cup \{C[t\sigma], s \approx t\}$
if $s\sigma \succ t\sigma$ and $C[s\sigma] \succ (s \approx t)\sigma$

Condensement: $N \cup \{C\} \triangleright_{\mathcal{S}} N \cup \{D\}$
if D is a condensement of C

Deduction: $N \triangleright_{\mathcal{S}} N \cup \{C\}$
if C is conclusion of an inference in \mathcal{S} from N

THEOREM 6.3 (BACHMAIR, GANZINGER 1994) Let $N_0 \triangleright_{\mathcal{S}} N_1 \triangleright_{\mathcal{S}} \dots$ be a **fair** superposition derivation, meaning that all conclusions of inferences in \mathcal{S} from “persistent” clauses are contained in $N \cup \text{Red}(N)$, where $N = \bigcup_i N_i$. Then N_0 is unsatisfiable if, and only if, $\perp \in N_i$ for some i .

Application to Lists

We analyse the clauses that can be derived from any set N containing the list equations L , together with finitely many flat (cf. Section 3.3) equations in car , cdr , cons , and constants, and with disequations $a \not\approx b$ between constants.

Choose \succ such that subterms are smaller than superterms and such that constants are smaller than any non-constant term. We observe:

- L is saturated.
- Disequations $a \not\approx b$ can only be reduced with equations between constants, and by reflexivity resolution.
- Every positive superposition inference produces an equation of depth ≤ 2 , and containing at most one cons . Example (a and b are constants):

$$\frac{\text{car}(a) \approx b \quad \text{cons}(\text{car}(x), \text{cdr}(x)) \approx x}{\text{cons}(b, \text{cdr}(a)) \approx a}$$

Decidability of List Constraints

12

THEOREM 6.4 Clausal validity in L can be decided in quadratic time.

Proof. Only a quadratic number (in the number of constants in the negated and flattened problem clause) of different consequences can be derived, each in time $O(1)$ \square

NB: It takes some work (dynamic programming, hash-cons) to make inferences computable in $O(1)$ each.

Application to Arrays

13

Axioms AR : (again implicitly universally quantified; sorted variables)

$$\begin{aligned} \text{read}(\text{write}(a, i, v), j) &\approx \text{read}(a, j) \vee i \approx j \\ \text{read}(\text{write}(a, i, v), i) &\approx v \end{aligned}$$

Extensionality Ext :

$$\forall a, a' (\forall i (\text{read}(a, i) \approx \text{read}(a', i)) \rightarrow a \approx a')$$

Problem: Decide clausal validity in $AR[+Ext]$.

Preprocessing: Negate and skolemize the problem clause. Flatten the resulting set of equations and disequations so that all equations become flat and only disequations between constants remain.

Specializing Superposition to AR

14

Choose \succ such that terms properly embedded into a term are smaller than the embedding term.^a Then, $\text{read}(\text{write}(a, i, v), j) \succ \text{read}(a, j)$, for any subterms a, i, j and v , as the second term is properly embedded in the first. (Delete $\text{write}, i,$ and v in the first term.) Also, subterms are then smaller in \succ than proper superterms.

As in the list example, selection is void, as there are no negative literals in AR and as all the other clauses, resulting from negating the problem clause, are unit literals.

We observe, that AR is saturated up to redundancy. In fact, the only inference possible is

$$\frac{\text{read}(\text{write}(a, i, v), i) \approx v \quad \text{read}(\text{write}(a', i', v'), j') \approx \text{read}(a', j') \vee i' \approx j'}{v \approx \text{read}(a, i) \vee i \approx j}$$

yielding a tautology.

^aAn embedded term is obtained by deleting any number of nodes, thereby moving the subtrees of a deleted node v (in any order) below the parent node of v .

Consequences from AR and Flat Equations

15

Adding any number of flat equations, the only consequences that can be derived are either (i) again flat equations; or else are (ii) clauses of the form

$$\text{read}(b, j_0) \approx t \vee i_1 \approx j_1 \vee \dots \vee i_m \approx j_m,$$

where t is either a constant or a term of the form $\text{read}(b, j)$, with j a constant, and with all other names denoting constants; or else are (iii) clauses of the form

$$\text{read}(a, J) \approx \text{read}(a', J) \vee i_1 \approx j_1 \vee \dots \vee i_m \approx j_m,$$

with a and the i_l constants, with J a variable, and the j_l either constants or the variable J .

If, in (ii) or (iii), $m > n^2 + n$, where n is the number of (index) constants appearing in the initial flat equations, the clause is either a tautology or can be condensed into a smaller clause. (One equation between indexes must be trivial or occur at least twice.)

Decidability of Array Constraints

16

THEOREM 6.5 Clausal validity in AR can be decided using \triangleright_S in exponential time.

Proof. Only an exponential number (in the number of constants in the negated and flattened problem clause C) of different consequences can be derived from $AR \cup \text{flatten}(\neg C)$, each in time $O(1)$. \square

As was shown in [Stump et al, LICS2001], the problem is NP-hard, so that this superposition-based decision procedure is probably optimal.

Eta Expansion

17

To extend the method to arrays with extensionality we first eliminate array disequations using the extensionality axiom:

Iteratively replace any disequation $a \not\approx b$ between array constants a and b by the 3 literals

$$\text{read}(a, k) \approx v, \text{read}(b, k) \approx w, v \not\approx w$$

with fresh constants k, v and w .

This terminates as the dimensionality of the constants v and w in the new disequation becomes smaller. That is, if a and b are of sort $\text{array}[n]$, then v and w are of sort $\text{array}[n-1]$, with $\text{array}[0]$ the sort of primitive (non-array) elements.

One might call this transformation η -expansion.

THEOREM 6.6 Let E be any set of ground equations and disequations over read, write and constants. Let E_η be the result of η -expanding E . Then E is satisfiable in a $AR + Ext$ -model if, and only if, E_η is satisfiable in a AR -model.

Proof. Exercise \square

THEOREM 6.7 Clausal validity in $AR + Ext$ can be decided using \triangleright_S in exponential time.

THEOREM 6.8 (STUMP ET AL, LICS 2001) The full first-order theory of $AR + Ext$ is undecidable

7 Combination of Decision Procedures

References

- [1] Baader, F. and Tinelli, C. (1997). A new approach for combining decision procedures for the word problem, and its connection to the nelson-open combination method. In W. McCune Ed., *Automated Deduction – CADE-14, 14th International Conference on Automated Deduction*, LNAI 1249 (Townsville, North Queensland, Australia, July 13–17, 1997), pp. 19–33. Springer-Verlag.
- [2] Barrett, C., Dill, D., and Stump, A. (2002). A generalization of Shostak’s method for combining decision procedures. In *Proc. FroCos 2002* (2002). LNCS, Springer-Verlag.

- [3] Bjørner, N. (1998). *Integrating decision procedures for temporal verification*. Ph. D. thesis, Stanford University.
- [4] Ganzinger, H. (2002). Shostak Light. In *Proc. CADE 2002*. LNCS, Springer-Verlag.
- [5] Kapur, D. (2002). A rewrite rule based framework for combining decision procedures. In *Proc. FroCos 2002* (2002). LNCS, Springer-Verlag.
- [6] Nelson, G. and Oppen, D. C. (1979). Simplification by cooperating decision procedures'. *ACM Transactions on Programming Languages and Systems* 2(2), pp. 245–257.
- [7] Rueß, H. and Shankar, N. (2001). Deconstructing Shostak. In *Proceedings of the Sixteenth IEEE Symposium On Logic In Computer Science (LICS'01)* (June 2001), pp. 19–28. IEEE Computer Society Press.

- [8] Shankar, N. and Rueß, H. (2002). Combining Shostak theories. In *Proc. RTA 2002*, Lecture Notes in Computer Science (2002). LNCS, Springer-Verlag.
- [9] Shostak, R. E. (1984). Deciding combinations of theories. *J. Association for Computing Machinery* 31(1) (January), 1–12.
- [10] Tinelli, C. and Harandi, M. (1996). A new correctness proof of the Nelson-Oppen combination procedure. In *1st Int'l Workshop on Frontiers of Combining Systems (FroCoS'96)*, Volume 3 of *Applied Logic Series* (1996). Kluwer Academic Publishers.

Reminder: Theories

19

In the syntactic or axiomatic view, a (first-order) theory is given by a set \mathcal{F} of (closed) first-order Σ -formulas, and then one is interested in the **models** $\text{Mod}(\mathcal{F})$ of \mathcal{F} , that is, the set of Σ -algebras satisfying \mathcal{F} :

$$\text{Mod}(\mathcal{F}) = \{\mathcal{A} \in \Sigma\text{-alg} \mid \mathcal{A} \models G, \text{ for all } G \text{ in } \mathcal{F}\}$$

Dually, in the semantic view, when given a class $\mathcal{T} \subseteq \Sigma\text{-alg}$ of structures **closed under isomorphism**, one is interested in the (first-order) **theory** $\text{Th}(\mathcal{T})$ of \mathcal{T}

$$\text{Th}(\mathcal{T}) = \{G \in F_{\Sigma}(X) \text{ closed} \mid \mathcal{T} \models G\}$$

which is the set of Σ -formulas that are satisfied in all structures \mathcal{A} in \mathcal{T} .

In the sequel we will assume the semantic view where a **theory** \mathcal{T} is a class of structures over a given signature Σ , closed under isomorphism, and not containing the the **trivial** one-element algebra. **That is, from now on we assume that algebras have at least 2 elements.**

Example: The **free theory** of a signature Φ of functions symbols (called **free**) is defined as $F^{\Phi} = \Phi\text{-alg}$, the entire class of (non-trivial) Φ -structures.

Combination of Theories and Models

20

Forgetting symbols. Suppose $\Sigma \subseteq \Sigma'$, that is, $\Omega \subseteq \Omega'$ and $\Pi \subseteq \Pi'$.

For $\mathcal{A} \in \Sigma'\text{-alg}$, by $\mathcal{A}|_{\Sigma}$ we denote the Σ -structure for which

$$\begin{aligned} U_{\mathcal{A}|_{\Sigma}} &= U_{\mathcal{A}} \\ f_{\mathcal{A}|_{\Sigma}} &= f_{\mathcal{A}}, \text{ for } f \text{ in } \Omega \\ p_{\mathcal{A}|_{\Sigma}} &= p_{\mathcal{A}}, \text{ for } p \text{ in } \Pi \end{aligned}$$

That is, we simply ignore the functions and predicates associated with symbols in $\Sigma' \setminus \Sigma$. $\mathcal{A}|_{\Sigma}$ is called the **restriction** of \mathcal{A} to Σ .

Amalgamating algebras. Let $\Sigma = \Sigma_1 + \Sigma_2$ (formed by uniting the respective symbol sets). A Σ -Algebra \mathcal{A} is called an **amalgamation** of Σ_i -algebras \mathcal{A}_i iff, for $i = 1, 2$, we have that $\mathcal{A}|_{\Sigma_i} = \mathcal{A}_i$.

On theories: Let $\Sigma \subseteq \Sigma'$, $\mathcal{T} \subseteq \Sigma'\text{-alg}$, $\mathcal{T}_i \subseteq \Sigma_i\text{-alg}$. We define:

$$\begin{aligned} \mathcal{T}|_{\Sigma} &= \{\mathcal{A}|_{\Sigma} \mid \mathcal{A} \in \mathcal{T}\} \\ \mathcal{T}_1 + \mathcal{T}_2 &= \{\mathcal{A} \in (\Sigma_1 + \Sigma_2)\text{-alg} \mid \mathcal{A}|_{\Sigma_i} \in \mathcal{T}_i, \text{ for } i = 1, 2\} \\ \mathcal{T}^{\Phi} &= \mathcal{T} + F^{\Phi} \end{aligned}$$

The problem: We assume theories \mathcal{T}_1 and \mathcal{T}_2 , respectively, and want to test whether E is satisfiable in $\mathcal{T}_1 + \mathcal{T}_2$ (that is, $\mathcal{T}_1 + \mathcal{T}_2 \not\models \forall \neg E$) where the E are constraints (multisets/conjunctions of equations and disequations, hence $\neg E$ a clause) over the combined theory $\mathcal{T}_1 + \mathcal{T}_2$.

Example: $\mathcal{T}_1 = \mathbb{L}$, $\mathcal{T}_2 = (\mathbb{Q}, +)$,
 $E = \{\text{car}(\text{cons}(x + 3, l)) \approx y + 1, y \not\approx x + 2\}$

Preprocessing: A constraint E_1, E_2 over $\mathcal{T}_1 + \mathcal{T}_2$ is called **pure** whenever the conjuncts E_j are \mathcal{T}_j -constraints, $j = 1, 2$. Purifying constraints is a matter of linear-time preprocessing. Example:
 $E = \{\text{car}(\text{cons}(x', l)) \approx y'\}, \{x' \approx x + 3, y' \approx y + 1, y \not\approx x + 2\}$

Assumptions: Satisfiability of constraints decidable for the \mathcal{T}_i ; signatures Σ_i of the theories disjoint

$\text{Sat}[\mathcal{T}_1] \times \text{Sat}[\mathcal{T}_2]$

$$E_1, E_2 \triangleright_{\text{NO}} \perp \quad \text{if } \mathcal{T}_i \models E_i \rightarrow \perp, \text{ for one } i$$

Branch $[\mathcal{T}_1, \mathcal{T}_2]$

$$E_1, E_2 \triangleright_{\text{NO}} E_1 \cup \{x \approx y\}, E_2 \cup \{x \approx y\} \quad | \quad E_1 \cup \{x \not\approx y\}, E_2 \cup \{x \not\approx y\}$$

whenever x and y are two different variables appearing in $E_1 \cup E_2$ such that $E_1 \cap E_2$ contains neither $x \approx y$ nor $x \not\approx y$.

We will be investigating soundness and completeness of several more or less restricted versions of this system. The restrictions will mainly concern the Branch rule which generates an exponential search space.

The inference rules $\triangleright_{\text{NO}}$ are to be applied don't-care non-deterministically (no backtracking) until termination.

“|” means non-deterministic (backtracking!) branching of the derivation into two subderivations. Derivations are, therefore, trees. All branches need to be reduced until termination.

Clearly, all derivation paths are finite since there are only finitely many **shared variables** in E_1 and E_2 , therefore the procedure represented by the rules is terminating.

We call a constraint configuration to which no rule applies **irreducible**.

THEOREM 7.1 (SOUNDNESS I) If all path in a derivation tree from E_1, E_2 end in \perp , then E_1, E_2 is unsatisfiable in $\mathcal{T}_1 + \mathcal{T}_2$.

Proof. Exercise. \square

THEOREM 7.2 (SOUNDNESS II) Let E'_1, E'_2 (different from \perp) be a configuration on a branch in a derivation from E_1, E_2 If E'_1, E'_2 is satisfiable in $\mathcal{T}_1 + \mathcal{T}_2$, so is E_1, E_2 .

Proof. Exercise. \square

Completeness

25

For completeness we need to show that if one branch in a derivation terminates with an irreducible configuration E_1, E_2 (different from \perp), then E_1, E_2 (and, thus, the initial constraint of the derivation) is satisfiable in the combined theory.

As E_1, E_2 is irreducible by **Sat**, the two constraints are satisfiable in their respective component theories, that is, we have two models $I_j \in \mathcal{T}_j$ satisfying E_j . We are left with combining the models into a single one that is both a model of the combined theory and of the combined constraint. These constructions are called **amalgamations**.

Amalgamation is not Always Possible

26

Amalgamation may fail because of cardinality reasons.

Suppose \mathcal{T}_1 are the Booleans, and \mathcal{T}_2 is the first-order theory of a three-element domain. In this case $\mathcal{T}_1 + \mathcal{T}_2$ is empty so that no constraint E can be satisfiable. On the other hand, the empty constraint E is trivially satisfiable both in \mathcal{T}_1 and \mathcal{T}_2 .

In the two amalgamation lemmas to follow, amalgamation will be possible because either one theory component has models of all cardinalities, or the two components have models of the same cardinality.

The Case of One Arbitrary and One Free Theory

27

The first case we shall treat is special in that one of the theories is free. That is, we have $\mathcal{T}_1 = \mathcal{T}$ arbitrary and $\mathcal{T}_2 = F^\Phi$ the free theory over a signature Φ of **free** function symbols.

We call a constraint **flat** if it contains only equations and disequations between variables, and equations of the form $f(x_1, \dots, x_n) \approx x$, with variables x_i as arguments of the function symbol f . The latter equations we call (function) **definitions** or (function) **rules** (for f).

A flat constraint D is called **unambiguous** if whenever $f(x_1, \dots, x_n) \approx x$ and $f(x'_1, \dots, x'_n) \approx x'$ are two different rules in D for the same function symbol f then either there exists an i such that $x_i \not\approx x'_i$ is in D , or else either $x = x'$ or $x \approx x'$ is in D .

A pair E, D of constraints is called **compatible** if any equation or disequation between two variables that appears in E also appears in D , and vice versa.

Amalgamation Lemma I

28

LEMMA 7.3 (AMALGAMATION LEMMA I) Let E be a constraint over a theory \mathcal{T} and D a constraint over the extension \mathcal{T}^Φ of \mathcal{T} with free function symbols Φ disjoint from the signature Σ of \mathcal{T} . Moreover assume that D is flat and unambiguous and that E, D is compatible. Then $E \wedge D$ is satisfiable in $\mathcal{T}^\Phi = \mathcal{T} + F^\Phi$ if and only if E is satisfiable in \mathcal{T} .

Let I be a model in \mathcal{T} and α an assignment of the variables in E such that $I, \alpha \models E$. We will extend the interpretation to the functions in Φ and the additional variables in D such that the interpretation also satisfies D . Since the pair E, D is compatible, I, α already satisfies the equations and disequations in D between variables. First we extend α to the additional variables in D in an arbitrary manner. Then we extend I by interpretations for the free function symbols as follows: If f is a free function symbol and $f(x_1, \dots, x_n) \approx x$ is a function definition in D , evaluate the x_i as well as x in I, α , yielding values a_i and c , respectively, and define $f_I(a_1, \dots, a_n)$ to be c . Define f_I arbitrarily at all other argument tuples of the domain of I . We have to show that f_I is well-defined. A potential ambiguity may arise from the presence of another definition $f(x'_1, \dots, x'_n) \approx x'$ for f in D . In that case, as D is unambiguous, either there exists an index j such that $x_j \not\approx x'_j$ is in D , or else either $x = x'$ or $x \approx x' \in D$. By compatibility of E and D these equations or disequations must be satisfied in I, α , so that no ambiguity can arise.

Nelson/Oppen $\mathcal{NCO}_{\mathcal{D}}[\mathcal{T}, \Phi]$ for Theory Extensions

Sat

$$E, D \triangleright_{\text{NO}} \perp \quad \text{if } \mathcal{T} \models E \rightarrow \perp$$

Compose

$$E, D \cup \{f(x_1, \dots, x_n) \approx x, f(y_1, \dots, y_n) \approx y\} \triangleright_{\text{NO}} \\ E \cup \{x \approx y\}, D \cup \{x \approx y, f(x_1, \dots, x_n) \approx x\}$$

if $E \cap D$ contains neither $x \approx y$ nor $x \not\approx y$, and if for each $1 \leq i \leq n$ either $x_i = y_i$ or else $x_i \approx y_i \in D$.

Branch-D

$$E, D \triangleright_{\text{NO}} E \cup \{x \approx y\}, D \cup \{x \approx y\} \quad | \quad E \cup \{x \not\approx y\}, D \cup \{x \not\approx y\}$$

if there are two rules $f(x_1, \dots, x_n) \approx z$ and $f(x'_1, \dots, x'_n) \approx z'$ in D such that

- (i) $x = x_i \neq x'_i = y$, for some $1 \leq i \leq n$
- (ii) for no index $1 \leq k \leq n$ the disequation $x_k \not\approx x'_k$ is in D , and
- (iii) $z \neq z'$ and $z \approx z' \notin D$.

Completeness

THEOREM 7.4 $\mathcal{NCO}_{\mathcal{D}}[\mathcal{T}, \Phi]$ is sound and complete for deciding satisfiability of constraints $E \wedge D$ over \mathcal{T}^{Φ} , provided E is a constraint over \mathcal{T} , D is a flat constraint over F^{Φ} and the pair E, D is compatible.

Proof. Both **Compose** and **Branch-D** preserve flatness of the D -part and the compatibility of the constraint pairs. Now suppose that E, D is irreducible by $\mathcal{NCO}_{\mathcal{D}}[\mathcal{T}, \Phi]$. That means that in particular E is \mathcal{T} -satisfiable. Irreducibility by **Compose** and **Branch-D** implies that D is unambiguous. We may now apply the amalgamation lemma 7.3 to infer that $E \wedge D$ is satisfiable in \mathcal{T}^{Φ} . \square

Observe that this result gives us another method for deciding the congruence closure problem, that is, the UWP in Φ -alg. Take \mathcal{T} to be \emptyset -alg, the theory of the empty signature.

Question: What would a constraint solver for \emptyset -alg look like? How efficient can it be implemented?

The Case of Stably Infinite Theories

A theory \mathcal{T} is called **stably infinite** if for any constraint E over \mathcal{T} whenever E is satisfiable in \mathcal{T} there exists a model $I \in \mathcal{T}$ of cardinality ω such that $I \models \exists E$.

Let E be a constraint over \mathcal{T} and let S be a set of variables of E . E is called **compatible** with an equivalence \sim on S if the constraint

$$E \wedge \bigwedge_{x \sim y} x \approx y \wedge \bigwedge_{x, y \in S, x \not\sim y} x \not\approx y \quad (1)$$

is \mathcal{T} -satisfiable whenever E is \mathcal{T} -satisfiable. This expresses that E does not contradict equalities between the variables in S as given by \sim .

Branching Implies Compatibility of Constraints

33

PROPOSITION 7.5 If E_1, E_2 is a pair of constraints over \mathcal{T}_1 and \mathcal{T}_2 , respectively, that is irreducible by $\text{Branch}[\mathcal{T}_1, \mathcal{T}_2]$ then both E_1 and E_2 are compatible with some equivalence \sim on the shared variables S of E_1 and E_2 .

Proof. Irreducible by the branching rule, for each pair of variables x and y , both E_1 and E_2 entail either $x \approx y$ or $x \not\approx y$. Choose \sim to be equivalence given by all (positive) variable equations between shared variables that are entailed by E_1 . \square

Amalgamation Lemma II

34

LEMMA 7.6 (AMALGAMATION LEMMA II) Let \mathcal{T}_1 and \mathcal{T}_2 be two signature-disjoint, stably infinite theories. Furthermore let E_1, E_2 be a pair of constraints over \mathcal{T}_1 and \mathcal{T}_2 , respectively, both compatible with some equivalence \sim on the shared variables of E_1 and E_2 . Then $E_1 \wedge E_2$ is satisfiable in $\mathcal{T}_1 + \mathcal{T}_2$ if, and only if, E_1 and E_2 are satisfiable in \mathcal{T}_1 and \mathcal{T}_2 , respectively.

Proof of the Amalgamation Lemma (“if”-Direction)

35

Assume that each of the E_j is \mathcal{T}_j -satisfiable, that is, there exist models I_j in \mathcal{T}_j and variable assignments α_j such that $I_j, \alpha_j \models E_j$. As the E_i are compatible with \sim , an equivalence on their shared variables, we may assume that the I_j and α_j are chosen to also satisfy the extended constraints in (1) with S the set of shared variables. Since the theories are stably infinite we may additionally assume that the I_j are of cardinality ω . Let ρ_j denote a bijection from \mathbb{N} to the domain of I_j . By construction whenever we have two shared variables x and y , $\alpha_1(x) = \alpha_1(y)$ if, and only if, $\alpha_2(x) = \alpha_2(y)$. Hence we may assume that whenever x is a variable occurring in both E_1 and E_2 that the bijections have been chosen such that $\rho_1^{-1}(\alpha_1(x)) = \rho_2^{-1}(\alpha_2(x))$. Now define I to be the structure having \mathbb{N} as its domain and such that if f is a symbol in any of the signatures Σ_j then $f_I(n_1, \dots, n_k) = \rho_j^{-1}(f_{I_j}(\rho_j(n_1), \dots, \rho_j(n_k)))$. Define $\alpha(x) = \rho_j^{-1}(\alpha_j(x))$ if x is a variable occurring in E_j . By construction of the ρ_j this definition is independent of the choice of j . Clearly $I|_{\Sigma_j}, \alpha \models E_j$, for $j = 1, 2$, hence $I, \alpha \models E_1 \wedge E_2$. Moreover, the $I|_{\Sigma_j}$ are isomorphic (via ρ_j) to I_j , and thus are models in \mathcal{T}_j , so that I is in $\mathcal{T}_1 + \mathcal{T}_2$ as required.

Branching on Shared Variables

36

Observe that for Proposition 7.5 and, hence, for the amalgamation lemma to hold only branching on shared variables is required:

$\text{Branch-S}[\mathcal{T}_1, \mathcal{T}_2]$

$$\frac{E_1, E_2}{E_1 \cup \{x \approx y\}, E_2 \cup \{x \approx y\} \quad | \quad E_1 \cup \{x \not\approx y\}, E_2 \cup \{x \not\approx y\}}$$

whenever x and y are two different shared variables in E_1, E_2 such that $E_1 \cap E_2$ contains neither $x \approx y$ nor $x \not\approx y$.

By

$$\mathcal{NO}_S[\mathcal{T}_1, \mathcal{T}_2] = (\text{Sat}[\mathcal{T}_1] \times \text{Sat}[\mathcal{T}_2]) + \text{Branch-S}[\mathcal{T}_1, \mathcal{T}_2]$$

we denote the Nelson-Oppen inference system instantiated with \mathcal{T}_1 and \mathcal{T}_2 , and where branching is restricted to the S-type.

THEOREM 7.7 $\mathcal{NO}_S[\mathcal{T}_1, \mathcal{T}_2]$ is terminating and complete for deciding satisfiability of pure constraints E_1, E_2 over $\mathcal{T}_1 + \mathcal{T}_2$ for signature-disjoint, stably infinite theories \mathcal{T}_1 and \mathcal{T}_2 .

Proof. Suppose that E_1, E_2 is irreducible by $\mathcal{NO}_S[\mathcal{T}_1, \mathcal{T}_2]$. Applying the amalgamation lemma 7.6 in combination with the proposition 7.5 we infer that E_1, E_2 is satisfiable in $\mathcal{T}_1 + \mathcal{T}_2$. \square

A theory \mathcal{T} is called **compact** if for each constraint E that is satisfiable in \mathcal{T} , either E is satisfied in some infinite model of \mathcal{T} , or else there exists a number m (depending on E) such that if $I \in \mathcal{T}$ is a finite model of E then $|I| \leq m$. (Since we do not consider theories having trivial models, the smallest such m will be always greater than 1.) Compactness means that if a constraint has finite models of unbounded cardinalities then it must also have an infinite model. First-order theories are compact, as are stably infinite theories.

A compact theory \mathcal{T} is called **convex** if for any finite set Γ of Σ -equations and for Σ -equations A_i , $1 \leq i \leq n$, whenever $\mathcal{T} \models \Gamma \rightarrow A_1 \vee \dots \vee A_n$, then there exists an index j such that $\mathcal{T} \models \Gamma \rightarrow A_j$. For convex theories, any clausal validity problem can be reduced to a linear number of validity problems for Horn clauses. As we only consider theories without trivial models convexity implies stable infiniteness.

THEOREM 7.8 If \mathcal{T} is convex then \mathcal{T} is stably infinite.

Proof. We shall prove the contrapositive of the statement. Suppose \mathcal{T} is compact but not stably infinite. Then there exists a satisfiable constraint E that has only finite models in \mathcal{T} . Let, respectively, E^+ and E^- denote the subset of positive and negative equations in E . As \mathcal{T} is compact all models of E are bound in cardinality by some number m . Now consider the clause $C = E^+ \rightarrow \neg E^- \vee \bigvee_i x \approx x_i$, with pairwise different fresh variables x and x_i , $1 \leq i \leq m$, not occurring in E . $\mathcal{T} \models C$, as the clause exactly expresses that all E models have size less than or equal to m . However, $\mathcal{T} \not\models E^+ \rightarrow e$, for any $e \in \neg E^-$ (as otherwise E would not be satisfiable), and also $\mathcal{T} \not\models E^+ \rightarrow x \approx x_i$, for each i , as otherwise \mathcal{T} would have trivial models which we have excluded. \square

[2] have proved this theorem for first-order theories.

EXAMPLE 7.9 Let \mathbb{Q} be the rational numbers with linear arithmetic and without the inequality predicates.^a In the signature of \mathbb{Q} we assume to have all rational numbers as constants, the binary addition operator $+$, and, for each rational number q , a unary operator $q \cdot$ multiplying its argument by q . An infinite structure, \mathbb{Q} is stably infinite. It is well-known that this theory is convex. In fact, suppose that $\mathbb{Q} \models \sum_i c_i x_i = 0 \vee \sum_i d_i x_i = 0$ and, for the purpose of deriving a contradiction, assume that there are tuples a_i and b_i of rational numbers witnessing non-convexity of this disjunction, that is, we have $\mathbb{Q} \models (\sum_i c_i a_i = 0 \wedge \sum_i d_i a_i \neq 0)$ and $\mathbb{Q} \models (\sum_i c_i b_i \neq 0 \wedge \sum_i d_i b_i = 0)$. But then $\mathbb{Q} \models (\sum_i c_i (a_i + b_i) \neq 0 \wedge \sum_i d_i (a_i + b_i) \neq 0)$, which is a contradiction. The general case of disjunctions with more than 2 linear equations is a bit more tricky. Allowing disequations is disjunctions in an easy extension.

^aWhen we have a specific structure such as \mathbb{Q} , we identify it with the theory that is the isomorphism class of this structure.

EXAMPLE 7.10 Let Φ be a signature. The class of all (non-trivial) algebras satisfying the empty set of first-order axioms, F^Φ (the free theory over Φ) trivially is a first-order theory and **closed under products**. Any such theory is convex (cf. below).

Hereby \mathcal{T} is called **closed under products**, if whenever A and B are models in \mathcal{T} then also their product $A \times B$ is in \mathcal{T} . $A \times B$ has as domain the cartesian product of the domains of A and B , and functions are defined component-wise as

$$f_{A \times B}((a_1, b_1), \dots, (a_n, b_n)) = ((f_A(a_1, \dots, a_n), f_B(b_1, \dots, b_n))).$$

THEOREM 7.11 If \mathcal{T} is closed under products, then \mathcal{T} is convex.

Proof. Exercise \square

LEMMA 7.12 Suppose \mathcal{T} is convex, E a \mathcal{T} -constraint, and S a subset of its variables. Let, for any pair of variables x and y in S , $x \sim y$ if, and only if, $\mathcal{T} \models E \rightarrow x \approx y$. Then E is compatible with \sim .

We show that with this choice of \sim the constraint (1) is satisfiable in \mathcal{T} whenever E is. Suppose, to the contrary, that E is satisfiable but (1) is not, that is,

$$\mathcal{T} \models E \rightarrow \bigvee_{x \sim y} x \not\approx y \vee \bigvee_{x, y \in S, x \not\sim y} x \approx y$$

or, equivalently,

$$\mathcal{T} \models E^+ \wedge \bigwedge_{x \sim y} x \approx y \rightarrow \neg E^- \vee \bigvee_{x, y \in S, x \not\sim y} x \approx y.$$

By convexity of \mathcal{T} , the antecedent implies one of the equations of the succedent. Since the equations $x \approx y$, with $x \sim y$, are entailed by E and since E is satisfiable this means that this equation must come from the last disjunct and it must be implied already by E^+ . In other words, there exists a pair of different variables x' and y' in S such that $x' \not\sim y'$ and $\mathcal{T} \models E \rightarrow x' \approx y'$ which is impossible.

Sat $[\mathcal{T}_1] \times$ Sat $[\mathcal{T}_2]$

$$E_1, E_2 \triangleright_{\text{NO}} \perp \quad \text{if } \mathcal{T}_1 \models E_1 \rightarrow \perp \text{ or } \mathcal{T}_2 \models E_2 \rightarrow \perp$$

Propagate $[\mathcal{T}_1, \mathcal{T}_2]$

$$E_1, E_2 \triangleright_{\text{NO}} E_1 \cup \{x \approx y\}, E_2 \cup \{x \approx y\}$$

if x and y are two shared variables in E_1, E_2 such that $\mathcal{T}_1 \models E_1 \rightarrow x \approx y$ and $\mathcal{T}_2 \not\models E_2 \rightarrow x \approx y$, or else $\mathcal{T}_2 \models E_2 \rightarrow x \approx y$ and $\mathcal{T}_1 \not\models E_1 \rightarrow x \approx y$.

THEOREM 7.13 If $\mathcal{T}_1, \mathcal{T}_2$ are two convex, signature-disjoint theories, the calculus $\mathcal{N}\mathcal{O}_{\mathcal{P}}[\mathcal{T}_1, \mathcal{T}_2]$ is sound and complete for satisfiability of pure configurations E_1, E_2 over $\mathcal{T}_1 + \mathcal{T}_2$.

Proof. We show that any unsatisfiable (in $\mathcal{T}_1 + \mathcal{T}_2$) configuration E_1, E_2 reducible by **Branch-S** in $\mathcal{N}\mathcal{O}_{\mathcal{S}}$ and irreducible by **Sat** is reducible by **Propagate** in $\mathcal{N}\mathcal{O}_{\mathcal{P}}$. Suppose $E_1 \wedge E_2$ is reducible by some **Branch-S**-branching. As convexity implies stable infiniteness, since irreducibility by **Sat** means that both projections E_j are satisfiable, we may apply the (contra-positive of the) amalgamation lemma 7.6 to infer that E_1, E_2 cannot be compatible with any partitioning of their shared variables S . If **Propagate** were not applicable then E_1 and E_2 would entail the same set E of equations between their shared variables. Therefore they would both be compatible with E (considered as an equivalence relation), which is a contradiction. \square

THEOREM 7.14 The Nelson/Oppen procedure $\mathcal{N}\mathcal{O}_{\mathcal{P}}[\mathcal{T}_1, \mathcal{T}_2]$ for convex theories requires $O(n^2)$ calls to the constraint solvers of the component theories.

Proof. There are at most $O(n^2)$ pairs of shared variables for which entailment in a component theory must be tested in order to check whether an equality needs to be propagated. \square

The combination procedures can be iterated to work with more than 2 component theories by virtue of the following observations where signature disjointness is assumed:

THEOREM 7.15 If \mathcal{T}_1 and \mathcal{T}_2 are stably infinite, so is $\mathcal{T}_1 + \mathcal{T}_2$.

Proof. $\mathcal{N}\mathcal{O}_{\mathcal{S}}[\mathcal{T}_1, \mathcal{T}_2]$ is a complete constraint checker for $\mathcal{T}_1 + \mathcal{T}_2$, so that if a constraint E over $\Sigma_1 + \Sigma_2$ is satisfiable, in any $\mathcal{N}\mathcal{O}_{\mathcal{S}}$ derivation from the purified form of E there exists a branch leading to some irreducible constraint E_1, E_2 entailing E . The amalgamation lemma 7.6 constructs a model of cardinality ω for E from the models of E_1 and E_2 . \square

THEOREM 7.16 If \mathcal{T}_1 and \mathcal{T}_2 are convex, so is $\mathcal{T}_1 + \mathcal{T}_2$.

Proof. By a similar argument exploiting the completeness of $\mathcal{N}\mathcal{O}_{\mathcal{P}}$. By convexity of the \mathcal{T}_i , the disequations in E_1, E_2 can be completely ignored for the **Propagate** rule, whenever both E_i are satisfiable in \mathcal{T}_i . \square

Shostak's procedure assumes the presence of a (unitary) unification algorithm for any of the built-in theories \mathcal{T} . More specifically it is assumed that there exists an effectively computable function solve such that, for any \mathcal{T} -equation $s \approx t$:

- (A) $\text{solve}(s \approx t) = \perp$ if, and only if, $\mathcal{T} \models s \not\approx t$;
- (B) $\text{solve}(s \approx t) = \emptyset$ if, and only if, $\mathcal{T} \models s \approx t$; and otherwise
- (C) $\text{solve}(s \approx t) = \{x_1 \Rightarrow u_1, \dots, x_n \Rightarrow u_n\}$ is a finite set of rewrite rules over Σ such that
 - (i) the x_i are pairwise different variables occurring in $s \approx t$;
 - (ii) the x_i do not occur in the u_j ; and
 - (iii) $\mathcal{T} \models \forall X[(s \approx t) \leftrightarrow \exists Y(x_1 \approx u_1 \wedge \dots \wedge x_n \approx u_n)]$, where Y is the set of variables occurring in one of the u_j but not in $s \approx t$, and $X \cap Y = \emptyset$.

If a function solve with these properties exists we call \mathcal{T} **solvable**.

$\text{solve}(s \approx t)$, if different from \perp , may be viewed as a substitution $\sigma = [u_1/x_1, \dots, u_n/x_n]$ (possibly the identity), written as a set of rewrite rules $\{x_1 \Rightarrow u_1, \dots, x_n \Rightarrow u_n\}$, that solves the \mathcal{T} -equation $s \approx t$. In the terminology of unification theory, σ is a unifier of $s \approx t$. In fact, by (iii), $\mathcal{T} \models s\sigma \approx t\sigma$. Conversely, if τ is a unifier of $s \approx t$, (iii) implies that $\mathcal{T} \models \exists Y(x_1\tau \approx u_1\tau \wedge \dots \wedge x_n\tau \approx u_n\tau)$. In case \mathcal{T} is an equational (Horn-) theory, the existing Y can be expressed as a substitution ρ with a domain included in Y . Then, $\mathcal{T} \models \tau(x) \approx \rho(\tau(\sigma(x)))$, for each variable in $s \approx t$. In other words, σ is a most general unifier of $s \approx t$.

Solutions can be parameterized by new variables, those in Y . It is assumed that in each calling context for solve , the variables in Y are fresh. Where this needs to be formalized we shall write $\text{solve}_Z(s \approx t) = S$, assuming that then the extra variables appearing in S are not in Z .

Examples of Solvable Theories

EXAMPLE 7.17 Consider again \mathbb{Q} , the rational numbers with linear arithmetic and without the inequality predicates. This theory is convex. A solver is obtained by isolating one of the variables in an equation.

EXAMPLE 7.18 Lists \mathbb{L} are a solvable theory.

EXAMPLE 7.19 Let $\mathbb{Z}/(3)$ be the theory of the three-element field obtained by considering the remainders from division by 3. Let the signature consist of the constants 0 and 1, and the binary addition $+$. Clearly, $\mathbb{Z}/(3)$ is solvable. For example, $a + a + 1 = b + b$ is solved by $a \Rightarrow 1 + b$. However $\mathbb{Z}/(3)$ is **not convex** as witnessed by the disjunction $x \approx 0 \vee x \approx 1 \vee x \approx 1 + 1$.

Shostak Canonizers

Normally, and in particular in Shostak's own paper Shostak-84, the method is presented such that in addition to a solver also a canonizer is required for the theory. Canonizers simplify terms by normalization and one may decide the word problem for the theory simply by checking identity of normal forms. Here, for methodological reasons we are only dealing with the solver and will introduce canonizers only later when we present more refined versions of the procedure. Therefore we also require solvers to satisfy property (B) which otherwise the canonizer takes care of. Note, however, that the "only if" part of (B) is implied by (C) anyway.

Constraint Solving $\mathcal{S}[\mathcal{T}]$ à la Shostak

Contradiction

$$U \cup \{s \approx t\}, R \triangleright_S \perp \quad \text{if } \text{solve}(s \approx t) = \perp$$

$$U \cup \{s \not\approx t\}, R \triangleright_S \perp \quad \text{if } \text{solve}(s \approx t) = \emptyset$$

Solve

$$U \cup \{s \approx t\}, R \triangleright_S U, R \cup S$$

where

- (i) $S = \text{solve}_X(s \approx t) \neq \perp$, with X the set of variables appearing in the antecedent,
- (ii) both s and t are irreducible by R .

Reduce

$$U[x], R \cup \{x \Rightarrow t\} \triangleright_S U[t], R \cup \{x \Rightarrow t\}$$

The inference system \mathcal{S} models the main idea in Shostak's procedure of how to employ the given solver for constraint simplification. As we shall prove in detail below, $\mathcal{S}[\mathcal{T}]$ refines $\text{Sat}[\mathcal{T}]$. The refinement involves the representation of constraints over \mathcal{T} by the pair of two constraints U and R . Here U contains the disequations and the “unsolved” positive equations, whereas R is a positive constraint in [solved form](#), a substitution derived from previous constraint solving steps. We assume that R is empty initially so that all constraint equations are unsolved.

The Contradiction rule solves a single, previously unsolved constraint equation. If the solver returns \perp for an equation or \emptyset for a disequation, the constraint is unsatisfiable. Those instances of $\text{Sat}[\mathcal{T}]$ that are not dealt with by Contradiction can be reduced by instances of Solve or Reduce (cf. Theorem 7.25), so that \mathcal{S} is in fact complete.

Solve solves Σ -equations. Soundness of this rule is a consequence of the soundness of the solver, cf. Proposition 7.21 below. More specifically, we only solve normalized equations in which both sides are irreducible by R . The reduce inferences are designed to compute those normal forms. The solved equation is deleted from U and its solution S is added to the solved form R . The rules added to R upon Solve are all of the form $x \Rightarrow w$, and are called [variable definitions](#). By Propositions 7.20 and 7.22, R always contains at most one definition for a variable and is terminating. Sets of constraints R with these properties we call [solved forms](#).^a Reduce expands variables in U by their definitions. Reducing equations before solving them is essential for keeping [variable definitions in solved forms unique](#).

^aIn the context of CLP when one speaks of solved forms one often requires that variables occurring on the right side of the rules are themselves irreducible. This is not the case for our notion of solved forms. Terminating, confluent rewrite systems, they do represent substitutions but not necessarily idempotent ones. Without introducing auxiliary variables for denoting (shared) subterms, we may get an exponential blow-up of the terms at the right sides of the rules if we always keep them in normal form.

PROPOSITION 7.20 Any rule set R appearing in an \mathcal{S} -derivation contains at most one definition for any variable.

Proof. The property is trivially true initially where R is empty. When adding a rule set S to R in Solve, if R contains a definition $x' \Rightarrow t'$, S cannot contain a rule for x' . Otherwise x' would have to occur in $s \approx t$, and the equation being solved at this step would not be irreducible with respect to R . \square

PROPOSITION 7.21 The inference system is sound. More specifically, (i) whenever $U, R \vdash_{\mathcal{S}} U', R'$ then $\mathcal{T} \models \exists X(U \wedge R) \rightarrow \exists X, Y(U' \wedge R')$ and $\mathcal{T} \models \forall X, Y(U' \wedge R' \rightarrow U \wedge R)$, with Y the variables in U', R' but not in U, R ; and (ii) if $U, R \vdash_{\mathcal{S}} \perp$ then $U \cup R$ is unsatisfiable in \mathcal{T} .

Proof. The only slightly interesting case is Solve. By the soundness properties of solve, we have that $\mathcal{T} \models \forall X[(s \approx t) \leftrightarrow \exists Y S]$, where Y are the new variables in $\text{solve}_X(s \approx t)$. The two implications to be shown for establishing (i) follow by simple logical calculations. \square

In the rewrite systems R , variables are considered as constants which can not be substituted by other terms. In this sense the systems R induce terminating rewrite relations.

PROPOSITION 7.22 If $U, R \vdash_S U', R'$ and if R is terminating so is R' .

Proof. Let us, for a configuration U, R with variables in X , define $x \succ^X y$ if, and only if, y occurs on the right side of a definition for x in R . R is terminating if, and only if, \succ^X is a well-founded partial ordering on X . (For the “if” part, use a lexicographic path ordering over some precedence $>^X$ for which $\Phi >^X X >^X \Sigma$, and which coincides with \succ^X on X to show termination of R .)

We now show that if \succ^X is a well-founded partial ordering on X and if $U, R \vdash_S U', R'$ then $\succ^{X'}$ is a well-founded partial ordering on X' , the set of variables in the new configuration. The only non-trivial case is when the derivation is by Solve where the new variable definitions S are added to R . However only equations $s \approx t$ irreducible by R are solved, so that no variable

appearing in s or t is reducible by R . Therefore any variable occurring on the right side of a rule in S is irreducible by R . Also, according to the definition of a solver, right sides of rules in S are irreducible by S . Consequently, $\succ^{X'}$ is well-founded. \square

PROPOSITION 7.23 The inference system \mathcal{S} is terminating.

Proof. We need to describe a well-founded ordering \succ on configurations with terminating rewrite systems R for which all inference rules are strictly monotone. Define \succ such that \perp is minimal. Moreover if $\kappa = U, R$ and $\kappa' = U', R'$ are two configurations with X and X' , respectively, the set of variables occurring in κ and κ' , let $\kappa \succ \kappa'$ whenever

- (i) $|U| > |U'|$; or else
- (ii) $|U| = |U'|$, $R = R'$, and $U \Rightarrow_R U'$.

This ordering is well-founded. For if in a sequence $\kappa_0 \succ \kappa_1 \succ \dots$ no equations are deleted from U no new rules can be introduced, and therefore $R_i = R_{i+1}$. As the rewrite relations in configurations are all terminating any such sequence must be terminating. Clearly, the rules in \mathcal{S} are strictly decreasing with respect to \succ . \square

The proposition in particular shows that the number of new variables introduced during a derivation must be finite, irrespective of the way a solver introduces them.

PROPOSITION 7.24 Let R be a solved form.

- (i) If F is a formula over \mathcal{T} in which all variables are irreducible by R then $\mathcal{T} \models R \rightarrow F$ if, and only if, $\mathcal{T} \models F$.
- (ii) Let \mathcal{T} be convex. If U is a set of Σ -disequations each of which is satisfiable in \mathcal{T} and irreducible by R , then $U \cup R$ is satisfiable in \mathcal{T} .

Proof. For (i) we observe that since R is a solved form it is satisfiable and logically equivalent to the substitution σ sending each variable x to its normal form with respect to R . The domain of σ does not contain any of the variables in F for otherwise F would not be irreducible. Therefore $\mathcal{T} \models R \rightarrow F$ iff $\mathcal{T} \models F\sigma$ iff $\mathcal{T} \models F$.

For (ii) first note that $\forall X(U, R \rightarrow \perp)$ is equivalent to $\forall X(R \rightarrow \neg U)$, where $\neg U$ is a disjunction of (positive) equations. Therefore, by convexity of \mathcal{T} , if $\mathcal{T} \models U, R \rightarrow \perp$ then $\mathcal{T} \models R \rightarrow u \approx v$ for some disequation $u \not\approx v$ in U . Now use (i) to infer that $\mathcal{T} \models u \approx v$ contradicting the satisfiability of $u \not\approx v$. \square

THEOREM 7.25 Let \mathcal{T} be a convex, solvable theory. If U, R is a terminal configuration of \mathcal{S} then $U \cup R$ is satisfiable in \mathcal{T} .

Proof. If no inference in \mathcal{S} can be applied to U, R then (i) U contains only negative equations, (ii) any of the disequations in U is satisfiable in \mathcal{T} , (iii) R is a solved form (cf. propositions 7.20 and 7.22), and (iv) any term appearing in U is irreducible by R . Now apply (ii) in Proposition 7.24. \square

If we are only interested in deciding the UWP for a solvable \mathcal{T} , \mathcal{T} need not be convex, as stated by this result:

THEOREM 7.26 Let \mathcal{T} be a solvable theory. If U, R is a terminal configuration of \mathcal{S} with $|U| \leq 1$, then $U \cup R$ is satisfiable in \mathcal{T} .

Proof. If no inference in \mathcal{S} can be applied to U, R then (i) U contains at most one negative equation $s \approx t$ and $\mathcal{T} \not\models s \approx t$; (ii) R is a solved form (cf. propositions 7.20 and 7.22); and (iii) any term appearing in U is irreducible by R . Now apply (i) in Proposition 7.24 to infer that $\mathcal{T} \models (R \rightarrow \neg U)$. \square